

REMARKS/ARGUMENTS

This Amendment is responsive to the Office Action mailed on August 11, 2008. Before entry of this Amendment, claims 1-3, 5-8, and 10-24 were pending. In this Amendment, claims 1, 7, 10, and 13 have been amended, claims 23-24 have been canceled, and new claims 25-26 are presented. After entry of this Amendment, which is respectfully requested, claims 1-3, 5-8, 10-22, and 25-26 will be pending.

I. Rejections under 35 U.S.C. §112

All pending claims (claims 1-3, 5-8, and 10-24) were rejected under 35 U.S.C. § 112 ¶ 1 as failing to provide support for "locat[ing at least some of the replication storage volumes] outside the respective failure boundary for any of the types of storage failure" (independent claims 1 and 10) Applicant thanks the Examiner for explaining the Examiner's position on p. 15 of the Office Action.

Applicant has amended claims 1 and 10, from which all the rest of the pending claims depend, to remove the offending language. Applicant therefore respectfully submits that the claims as pending comply with the written description requirement, and requests that the rejections under §112 be withdrawn.

II. Rejections under 35 U.S.C. §103

Claims 1-3, 5-8, 10-20, and 23-24 were rejected under 35 U.S.C. § 103(a) as being unpatentable (obvious) over Bridge (US 6,530,035) in view of Iwami et al. (US 2002/0112030) and Ohran et al. (US 2002/0112134). The balance of the pending claims (claims 21-22) were rejected under 35 U.S.C. § 103(a) as being unpatentable (obvious) over Bridge in view of Iwami and Ohran, supported by Mehta et al. (US 2004/0250042). (Mehta is cited for the proposition that horizontal and vertical accessing of storage volumes were well known in the art at the time of the invention.) Applicants respectfully submit that the relied-upon references do not teach or suggest each element of the claims.

In particular, Iwami does not disclose "one of the plurality of failure boundaries being determined using error correction group information of the primary storage volumes and replication storage volumes" (emphasis added). Iwami is directed to determining logical disks which can accommodate a user's required speed and capacity using error correction group information (see Iwami paragraph [0032]) and not "failure boundaries." Thus, combining the teachings of Iwami with Bridge, which does not disclose "failure boundaries" being determined using error correction group information of the primary and replication storage volumes (see Office Action p. 5), would not result in the invention as claimed.

The independent claims 1 and 10 have been amended to recite "a full backup" and an "incremental backup," such that "a full backup to be stored has replication storage volumes assigned across each failure boundary" (emphasis added) and "an incremental backup is able to be stored having replication storage volume within at least one failure boundary" (emphasis added). Support for this is in the specification, for example in paragraphs [0055] and [0057].

As argued in Applicant's May 5, 2008 response, all of the full backup data would be lost in the sample embodiments taught by Ohran if a mass storage device upon which the full data backup resides failed. If the preservation memory storing the incremental backup data was in fact in a different failure boundary than the mass storage device failure, then the incremental backup data would still be available. Otherwise, if the preservation memory did reside in the same failure boundary as the mass storage device, then the incremental data would be fully lost as well. Ohran is effectively teaching that these different types of data, incremental backup data and full backup data, can be independently stored in different failure boundaries. The teaching in Ohran contrasts with what is claimed in the present invention. Ohran does not teach how to store a given type of content across failure boundaries. Storing data across a failure boundary affects what portions of the content is lost when a failure event occurs.

If the full backup data is stored as claimed, then only a subset of the full backup data would be lost for a failure condition used to define a failure boundary. In Ohran, all of the data could be lost due to a failure condition such as the mass storage device failing. If the failure of the mass storage device is used as a failure boundary as in the Applicants' claims, then only a

subset of the full backup data would be lost due to that failure event. Ohran teaches steps that can be taken to try to recover from this failure, but the steps taken after a failure event are not the subject matter of the Applicants' claims. The Applicants' claims deal with how data is stored before a failure event occurs.

Furthermore, the combination of Bridge, which discloses a mirroring technique across failure boundaries (*see, e.g.*, Bridge col. 2, lines 42-57), Iwami, which is directed to allocating disk resources for speed (*see, e.g.*, Iwami abstract), and Ohran, which is directed to restoring data from incremental backups (*see, e.g.*, Ohran paragraph [0012]), would not result in a method in which "a full backup to be stored has replication volumes assigned across each failure boundary . . . and . . . an incremental backup is able to be stored having replication storage within at least one failure boundary." The combination of the cited references apparently would result in the more conventional technique of storing a single type of backup across a failure boundary as opposed to storing "a full backup" and "an incremental backup" across and within failure boundaries. Additionally, the teachings of the other relied upon reference, Mehta, do not make up for the deficiencies of the other cited art. Mehta is directed to efficiently reading arbitrary patterns of address locations from a memory of a signal processor (*see, e.g.*, Mehta abstract). Applicants therefore request the withdrawal of the rejections of the independent claims and all claims depending thereon for at least the reasons stated above.

Claims 7 and 13 have been amended to recite that "the determining [of failure boundaries] includes using established reliabilities of the . . . volumes." This is supported in the specification, for example at paragraph [0040] and [0042] and in FIG. 6(e). If the Examiner objects to the phrase "established reliabilities" because it is not expressly stated in the application, then Applicant would be willing to put forth alternate language.

Using the established reliabilities of the volumes, whether entered by an administrator as illustrated in paragraph [0042] or queried from a group table as in FIG. 6(e), allows the empirical determination of failure boundaries of storage volumes, even if they use the same makes and models of disk drives. This could allow computer software to dynamically compensate for differences in reliability of different persistent storage units. For example,

RAID1 may be found to be more reliable than RAID5, so that a failure boundary may be automatically drawn around RAID5 for the types of failures to which RAID5 is more susceptible (e.g., failures caused by voltage surges, vibrations, hot/humid weather, high throughputs for long durations, striping/dithering by database logical volume managers, or just unknown causes). Bridge, Iwami, and Ohran neither teach or suggest using established reliabilities of the volumes to determine failure boundaries. At best, one could assume from the references that "two disk drives on a common controller" are reliable enough for certain applications and not as reliable for others (as alluded to in the Office Action, p. 11). However, this assumption falls far short of using "established reliabilities" to determine failure boundaries, even if an established reliability for a volume is simply summarized to the user as "High," as in FIG. 6(e). Thus, Applicants submit that claims 7 and 13 distinguish over the prior art and are patentable for at least the reasons stated above.

III. Amendment to the Claims

Unless otherwise specified or addressed in the remarks section, amendments to the claims are made for purposes of clarity, and are not intended to alter the scope of the claims or limit any equivalents thereof. The amendments are supported by the specification and do not add new matter.

IV. New claims

Presented are new dependent claims 25-26 which recite that "the incremental backup is stored in an ATA disk drive based storage system." Support for this is in the specification, for example at paragraph [0051]. Applicant submits that the combination of Bridge, Iwami, and Ohran does not teach or suggest that incremental backups are stored to an ATA disk drive based storage system, which incorporates high speed with low cost as opposed to traditional magnetic tape drives.

Appl. No. 10/767,247
Amdt. dated October 3, 2008
Amendment under 37 CFR 1.116 Expedited Procedure
Examining Group 2185

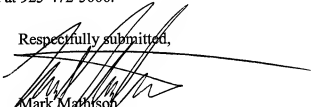
PATENT

CONCLUSION

In view of the foregoing, Applicants believe all claims now pending in this Application are in condition for allowance and an action to that end is respectfully requested.

If the Examiner believes a telephone conference would expedite prosecution of this application, please telephone the undersigned at 925-472-5000.

Respectfully submitted,



Mark Mathison
Reg. No. 57,556

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, Eighth Floor
San Francisco, California 94111-3834
Tel: 925-472-5000
Fax: 415-576-0300

MPM:sea
61542886 v1